Perthes Disease Classification Using Shape and Appearance Modelling

DOI:

[10.1007/978-3-030-11166-3_8](https://doi.org/10.1007/978-3-030-11166-3_8)

Document Version

Accepted author manuscript

[Link to publication record in Manchester Research Explorer](#)

Citation for published version (APA):

Davison, A., Cootes, T., Perry, D. C., Luo, W., & Lindner, C. (2019). Perthes Disease Classification Using Shape and Appearance Modelling. In *6th MICCAI Workshop on Computational Methods and Clinical Applications in Musculoskeletal Imaging* (Vol. 11404 LNCS, pp. 86-98). (Lecture Notes in Computer Science; Vol. 11404). Springer Nature. https://doi.org/10.1007/978-3-030-11166-3_8

Published in:

6th MICCAI Workshop on Computational Methods and Clinical Applications in Musculoskeletal Imaging

Citing this paper

Please note that where the full-text provided on Manchester Research Explorer is the Author Accepted Manuscript or Proof version this may differ from the final Published version. If citing, it is advised that you check and use the publisher's definitive version.

General rights

Copyright and moral rights for the publications made accessible in the Research Explorer are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

Takedown policy

If you believe that this document breaches copyright please refer to the University of Manchester's Takedown Procedures [<http://man.ac.uk/04Y6Bo>] or contact uml.scholarlycommunications@manchester.ac.uk providing relevant details, so we can investigate your claim.

Perthes Disease Classification Using Shape and Appearance Modelling

Adrian K. Davison¹(✉), Timothy F. Cootes¹, Daniel C. Perry^{2,3},
Weisang Luo³, Medical Student Annotation Collaborative^{2,3}, and
Claudia Lindner¹

¹ The University of Manchester, Centre for Imaging Sciences, UK

² University of Liverpool, UK

³ Alder Hey Children's Hospital, UK

`adrian.davison@manchester.ac.uk`

Abstract. We propose to use statistical shape and appearance modelling to classify the proximal femur in hip radiographs of children into Legg-Calvé-Perthes disease and healthy. Legg-Calvé-Perthes disease affects the femoral head with avascular necrosis, which causes large shape deformities during the growth-stage of the child. Further, the dead or dying bone of the femoral head is prominent visually in radiographic images, leading to a distinction between healthy bone and bone where necrosis is present. Currently, there is little to no research into analysing the shape and appearance of hips affected by Perthes disease from radiographic images. Our research demonstrates how the radiographic shape, texture and overall appearance of a proximal femur affected by Perthes disease differs and how this can be used for identifying cases with the disease. Moreover, we present a radiograph-based fully automatic Perthes classification system that achieves state-of-the-art results with an area under the receiver operator characteristic (ROC) curve of 98%.

Keywords: Computer-aided diagnosis, Perthes disease, random forests, radiographs, paediatrics, shape modelling, appearance modelling.

1 Introduction

Legg-Calvé-Perthes disease (Perthes) is an idiopathic disease in children between the ages of 2 – 14 years, with boys being affected 5 times more than girls [1]. The age-of-onset follows a lognormal distribution i.e. the disease has the tendency to affect younger rather than older children [2, 3]. Perthes disease is usually analysed through radiographic images in the anterior-posterior (AP) or frog lateral views of the hip. It is not yet known what exactly causes Perthes disease, however, environmental, congenital and socio-economic issues have been associated with Perthes [4]. There is also currently no defined best practice on how to treat the disease, and the decision is usually determined by the treating surgeon. One way of helping to identify and treat Perthes in clinical practice is to use classification methods. Three main categories exist: classification of the stage of disease

Fig. 1. An example of 58 points annotated on (a) a healthy child hip and (b) a hip affected by Perthes disease. Note that the landmark points were placed automatically using a Random Forest Regression-Voting system (see Section 3.3 for details).

progression [5, 6], classification of prognostic outcomes [7–9] and classification of the patient’s long term outcome [10].

There are only very few methods that utilise computer vision to analyse and study Perthes disease and, to the best of our knowledge, so far no computer vision based methods have been presented to classify between Perthes hips and healthy hips. A semi-automatic radiograph-based method was created for the quantitative analysis of the hips of children with Perthes [11], where manual landmark points initialised the femoral head contour, and a gradient operator with linear interpolation was used for the final contour location. The bone loss in the affected hip was identified by comparing the area included in the affected contour with that in the contour of the contra-lateral unaffected hip using the brightness of pixels (from 0 to 255 grey levels).

Chan et al. [12] used statistical shape modelling to understand the morphological deformities in both Perthes disease and slipped capital femoral epiphysis (SCFE) using 3D CT scans. Their results showed that the analysis of femoral shape during growth and in various disease stages are contributing to the understanding of normal and abnormal hip shape deviations, the latter of which may affect the risk of developing hip osteoarthritis.

Currently, in clinical practice, any method to diagnose or classify stages of Perthes disease or to determine patient outcomes are done manually by the treating surgeon.

In this study, we investigate how the radiographic shape, texture and appearance of children’s hips can be used to distinguish children’s hips affected by Perthes disease from healthy children’s hips. Our analysis is based on outlining the proximal femur with landmark points and applying statistical shape and appearance modelling [13, 14]. We test each of the three parameter sets (shape, texture, and appearance) individually to identify if any one of them outperforms the others as a classification feature. We use a Random Forest classifier (RF) [15] for this task, comparing our automatically obtained classification results to data manually categorised by clinicians (Perthes vs. healthy hips).

Further, we investigate the classification performance when the landmark point positions are obtained fully automatically via a Random Forest Regression-Voting (RFRV) [16, 17] system, rather than using manual landmark annotations (i.e. point positions). The latter are very time-consuming to obtain and prone to inconsistencies. Therefore, creating a fully automatic method to both annotate the hip and classify disease status would greatly reduce the amount of time clinicians need to spend analysing patient data, and facilitate the integration of such a system into the clinical workflow.

Finally, we analyse how classification results based on manual landmark annotations compare to classification results based on fully automatically obtained landmark annotations. Our results demonstrate that our fully automatic classification system is able to replicate the healthy vs Perthes classification by clinicians, with an area under the ROC curve (AUC) of 98%.

2 Background

Locating landmarks on medical images is an important first step in many musculoskeletal analysis tasks, particularly those requiring geometric measurements of the shape of structures (see Fig. 1 for a landmark annotation example). Many methods have been proposed for automating landmark localisation, with some of the most effective using Random Forest Regression-Voting (RFRV) [16, 17] which has been used for automatically locating landmarks along the proximal femur in radiographs of adult hips [16].

Techniques for analysing human skeletal structures [18] and their associated diseases are well established in describing the differences between healthy and diseased bone. Waarsing et al. [19] constructed statistical shape and appearance models for the left and right proximal femurs for cases of osteoarthritis. Their results show that subtle shape and appearance changes can be identified with these models in cases where traditional clinical measures might miss them. Whitmarsh et al. [20] used statistical shape and appearance models to distinguish fractured bones from a non-fractured control group using Fisher Linear Discriminant Analysis. They concluded that the proposed model-based fracture risk estimation method may improve upon the current standard in clinical practice.

Thomson et al. [21] analysed the shape and texture of the tibia in radiographs of osteoarthritis-affected knees using Random Forests for classification.

Table 1. A breakdown of the Perthes and healthy radiograph dataset with the total numbers of annotated and unannotated images.

	Annotated	Unannotated	Total
Healthy	1109	284	1393
Perthes	70	317	387
Total	1179	601	1780

Their fully automatic system achieved an AUC of 0.849 when combining both radiographic shape and texture, up from 0.789 when using shape alone. Their results demonstrate the effectiveness of using both radiographic shape and texture for classification.

Radiographic shape and appearance have also been used to estimate bone age [22] from radiographs of children’s hands using a RFRV system. The method achieved mean absolute prediction errors of 0.57 years and 0.58 years for females and males, respectively.

3 Method

3.1 Data Collection and Annotation

The dataset consists of (a) 387 AP pelvic radiographs of children (aged between 2 – 11 years) affected by Perthes and (b) 1393 radiographs of children not affected by Perthes (aged between 2 – 11 years). 1109 of the healthy cases, and 70 of the diseased cases were manually annotated with 58 points as shown in Fig. 1. There were no manual annotations for the remainder of the images. For the sake of convenience, the annotated dataset and unannotated dataset will henceforth be referred to by “Data-A” and “Data-U”, respectively. See Tab. 1 for a breakdown of the total number of radiographs.

This dataset is very challenging due to the natural growth stage during childhood, meaning the femur has growth areas such as the femoral head and greater trochanter. In addition, Perthes disease can have a significant effect on radiographic shape and appearance. Fig. 2 shows some examples of the challenging nature of the dataset. Even clinicians consider the task of manually annotating these hips (to create a ground truth) difficult, which increases the complexity of developing a system that would do this automatically.

3.2 Shape and Appearance Modelling

A statistical shape model (SSM) consists of a linear model of the distribution of a set of landmarks across a set of images. In the following we provide a brief summary on how to generate an SSM, for more details see [13]. To generate an SSM, the training data is a set of n images $\mathbf{I}$ with annotations $\mathbf{x}_l$ of a set of N landmark points $l = 1, \dots, N$ on each image. In this study, we use both

Fig. 2. Examples of healthy hips: (a) shows an older child but with visible growth plates on the femoral head and greater trochanter, and (b) is a 2 year old child with an early growth stage femoral head. Examples of Perthes hips: (c) demonstrates the difficulty of identifying the outline of the superior femoral head in some cases, and (d) gives an example of the extreme deformities of the femoral head.

manually obtained landmark positions and automatically obtained landmark positions. To begin, each image is aligned to a standard reference frame using a similarity transformation T with parameters θ . An SSM can then be created by applying principal component analysis (PCA) to all n training shapes in the reference frame, generating a linear model of shape variation that describes the position of each point l by

$$\mathbf{x}_l = T_\theta(\bar{\mathbf{x}}_l + \mathbf{P}_{sl}\mathbf{b}_s) \quad (1)$$

where $\bar{\mathbf{x}}_l$ is the mean position of the landmark point in the reference frame, $\mathbf{P}_{sl}$ is a set of modes of shape variation relating to the landmark point, and $\mathbf{b}_s$ are the shape model parameters.

Using dimensionality reduction, SSMs can be used to provide a compact quantitative description of the shape of the bone, which is very useful for classification tasks. However, SSMs only consider the distribution of the landmark point positions and hence only describe the radiographic shape of the bone. Perthes disease is known for avascular necrosis of the femoral head, which in radiographs shows as opposite pixel intensities compared to healthy bone. Statistical appearance models (SAMs), as used in the well-known Active Appearance Models [14] method, apply PCA-based linear modelling to both landmark point positions (i.e. shape) and pixel intensities (i.e. texture).

In the following we provide a brief summary on how to generate an SAM, for more details see [14]. To build a texture model, a patch comprising the set of landmark points is sampled from each training image. All patches are shape-normalised and texture-normalised to generate shape-free patches where global lighting variations have been removed. Each patch is then sampled into a texture vector $\mathbf{g}$ representing the texture of a particular training image in the reference frame. Given the set of n normalised texture vectors, PCA can be

applied to generate a linear texture model

$$\mathbf{g} = \bar{\mathbf{g}} + \mathbf{P}_g \mathbf{b}_g \quad \text{and} \quad \mathbf{b}_g = \mathbf{P}_g^T (\mathbf{g} - \bar{\mathbf{g}}) \quad (2)$$

where $\bar{\mathbf{g}}$ is the mean texture, $\mathbf{P}_g$ are the modes of texture variation, and $\mathbf{b}_g$ are the texture model parameters.

SAMs combine both shape and texture models to also capture correlations between shape and texture. Following the description above, the appearance of an image can be summarised using shape parameters $\mathbf{b}_s$ and texture parameters $\mathbf{b}_g$. To generate an SAM, appearance vector $\mathbf{b}$ can be defined by

$$\mathbf{b} = \begin{pmatrix} \mathbf{W}_s \mathbf{b}_s \\ \mathbf{b}_g \end{pmatrix} \quad (3)$$

where $\mathbf{W}_s$ is a diagonal matrix of weights to account for the difference in units between the shape and texture models (e.g. coordinates vs pixel intensities). Applying PCA to $\mathbf{b}$ yields an SAM given by

$$\mathbf{b} = \mathbf{P}_c \mathbf{c} \quad (4)$$

where $\mathbf{P}_c$ is a set of modes of appearance variation, and $\mathbf{c}$ are the appearance model parameters. Applying SSMs and SAMs to radiographic images provides a meaningful way to capture the variation in radiographic shape and texture that may allow to distinguish between proximal femurs affected by Perthes disease and healthy proximal femurs. In this study, we explore the effectiveness of using (i) shape model parameters $\mathbf{b}_s$; (ii) texture model parameters $\mathbf{b}_g$; or (iii) appearance model parameters $\mathbf{c}$ for classifying diseased and healthy hips.

3.3 Automatic Landmark Annotation

In light of applying the proposed technology in clinical practice it would be necessary for the system to be fully automatic. That is, the proposed classification system would need to be able to automatically place the 58 landmark points. For this purpose, we trained a RFRV system as presented in [16,17]. We used Data-A as training data for the system and performed five-fold cross-validation experiments (i.e. the data was randomly split into five even blocks and each block was used once for testing with the remaining blocks used for training). To be able to estimate the performance of a fully automatic classification system and compare this to a classification system based on manual ground truth, we combined the test results of all five folds to obtain a set of automatic annotations for Data-A. Note that because we used five-fold cross-validation experiments to generate the automatic annotations for Data-A, all automatic landmark point positions were obtained without training and testing on the same data. Comparing the manual and automatic landmark annotations for Data-A shows that the RFRV system achieved a point-to-curve-error of 4% of the femoral shaft width for 95% of all 1179 images and a median accuracy of less than 1.8% of the femoral shaft width.

Furthermore, the majority of our Perthes data (317 images) are unannotated. To allow us to utilise this data, Data-U, for evaluating the classification performance in this study, we randomly chose one of the five cross-validation RFRV systems trained on Data-A and used this to fully automatically annotate all images in Data-U.

4 Evaluation

To classify between Perthes and healthy hips, we obtained the shape, texture and appearance model parameter values based on annotated (manually and/or automatically) proximal femurs (healthy and Perthes) as shown in Fig. 1. We used the shape, texture and appearance model parameter values as classification features. Throughout the classification evaluation, we performed 5-fold cross-validation experiments (i.e. the data was randomly split into five even blocks and each block was used once for testing with the remaining blocks used for training) and we report the average classification performance over all five runs.

For all classification experiments, we used Random Forests (RF) [15] with 500 trees as the classifier. We applied bootstrapping, and the number of features to consider for each node split was set to $\sqrt{n\_features}$ with $n\_features$ being the number of shape, texture or appearance model parameter values. When obtaining the shape, texture and appearance model parameter values, we constrained the number of modes of variation such that the texture model explained 85% and the shape/appearance models each explained 98% of the data variation. We report the results using receiver operator characteristic (ROC) curves that show the true positive rate (TPR) against the false positive rate (FPR), along with the area under the curve (AUC).

4.1 Data-A Perthes Classification

Data-A includes manual annotations for 70 Perthes and 1109 healthy images. The classification results based on the model parameters obtained from the manual annotations (see Fig. 3) show that texture does not perform as well as shape or appearance. The best classification results were obtained when using the shape or appearance model parameter values with an AUC of 0.93 (SD: ± 0.06) and 0.93 (SD: ± 0.03) respectively.

4.2 Balanced Data-A Perthes Classification

Data-A has an imbalance between classes (70 Perthes cases vs. 1109 healthy cases) which could be a disadvantage in the above experiments. To investigate the impact of this class imbalance on performance, we took a random subset of 100 healthy hips from Data-A such that the classes were much closer in number, and re-ran the classification experiments. Fig. 4 shows that this leads to improved classification results for all models, significantly boosting the texture model classification performance with an AUC of 0.96 (SD: ± 0.02).

Fig. 3. Cross-validation ROC curves for Perthes-healthy classification when using shape, texture or appearance parameters based on Data-A (70 Perthes and 1109 healthy). All results were obtained using manual ground truth landmark annotations.

Training the classifier on a proportionally large amount of normal, healthy hips can create a bias towards the radiographic shape and appearance of healthy hips. Due to the effects of disease, Perthes cases show a much wider variation in the radiographic shape, texture and appearance parameter values. It may, thus, be beneficial to keep the datasets as balanced as possible. The results in Fig. 4 demonstrate the potential performance improvements when using a balanced dataset.

4.3 Fully Automatic Shape and Appearance Analysis

Our fully automatic system uses RFRV to locate the landmark points without the need for any manual intervention. Fig. 5 shows the fully automatically obtained classification results for Data-U (284 healthy and 317 Perthes) where the model parameter values were obtained from the automatically located landmark points. The best performance was achieved when using the shape or appearance model parameters with an AUC of 0.98 (SD: ± 0.01). Overall, the classification results for Data-U (using automatic landmark annotations) are better than the results obtained for Data-A (using manual landmark annotations).

Similar to the manual annotation results, the shape and appearance parameters outperform the texture parameters in the fully automatic analysis. It is noteworthy that Data-U contains many more Perthes cases than Data-A. Therefore, this setting is a more challenging task due to the increased range of radiographic shape and appearance variations across Perthes cases. This is in

Fig. 4. Cross-validation ROC curves for Perthes-healthy classification when using shape, texture or appearance parameters based on a subset of Data-A with a more balanced number of healthy and Perthes data (70 Perthes and 100 healthy). All results were obtained using manual ground truth landmark annotations.

particular the case because the RFRV system used to automatically locate the landmark points in Data-U was trained using Data-A which only includes 70 Perthes cases in total.

4.4 Manual Versus Automatic Classification

The automatic classification results for Data-U show an improvement in performance over the manual classification results for Data-A. However, this improvement in performance may originate from the difference in datasets. To directly compare the fully automatic classification performance to a classification system based on manual landmark annotations, we re-ran the classification experiments for Data-A using the automatically obtained Data-A landmark annotations (see Section 3.3) rather than the manual ground truth Data-A annotations.

Fig. 6 gives the results of the comparison for each of the parameter sets. The results show that the fully automatic classification system performs better than the classification system based on manual landmark annotations. The best performance was obtained when using the appearance model parameters with AUCs of 0.96 (SD: ± 0.02) and 0.93 (SD: ± 0.03) for the automatic and manual systems, respectively. These results demonstrate that we are able to fully automatically annotate diseased and healthy hips, and accurately classify the data, even when the data is imbalanced.

Fig. 5. Cross-validation ROC curves for the fully automatically obtained classification results for Data-U when using shape, texture or appearance model parameters. All landmark point positions were obtained automatically using the developed RFRV system.

5 Discussion and Conclusions

We have evaluated a radiograph-based classification system to distinguish proximal femurs affected by Perthes disease from healthy ones by using shape, texture and appearance model parameters. We have investigated how each set of parameters performs using a Random Forest classifier to identify healthy and Perthes hips. Our experiments show that the combination of shape and texture (appearance) performs best, achieving an AUC of 98% when using a fully automatic classification system.

In all our experiments, except for the balanced dataset experiments, classification based on shape model parameters outperformed the classification based on texture model parameters. Although the radiographic texture of the proximal femur may be affected by the radiolucency effect (caused by the dying bone of the femoral head in the early-mid stages of Perthes [9], changes in bone shape seem to be more discriminative. This highlights the impact of Perthes disease on the (radiographic) shape of the proximal femur. However, the discriminatory power of texture may improve when using a balanced dataset.

Our comparison of the performance of a fully automatic classification system to a classification system based on manual landmark annotations demonstrates that improved performance can be achieved when using automatically identified landmark positions. A possible explanation for this is that the automatic an-

Fig. 6. Cross-validation ROC curves for the comparison between the manual and automatic classification results for Data-A when using (a) shape, (b) texture and (c) appearance model parameters.

notations are placed more consistently, reducing random errors introduced by manual landmark annotations.

We have shown a viable system based on statistical shape and appearance models to automatically classify whether a hip is affected by Perthes disease or not. The proposed system would save clinicians' time, and produce accurate and robust results in clinical practice. In addition, such a system would be of benefit to support less experienced clinicians' or in a non-specialty clinical setting.

Further work will add more manually annotated diseased data during training for the comparison of the agreement between clinical diagnosis (Perthes vs. healthy hips) and the outputs of the automatic system. As Perthes is a rare disease, the availability of Perthes data compared to healthy data is low. Future work will focus on utilising a balanced dataset with as many Perthes cases as possible for developing (i.e. training) an automatic classification system, and then evaluating the system on an unseen imbalanced dataset to reflect the data availability in clinical practice.

Moreover, the system could be extended to use radiographic shape and appearance in combination with clinical data to also classify (i) the stage of disease

progression [5,6]; (ii) prognostic outcomes [7–9]; and (iii) long term patients’ outcomes [10]. Once we have collected more data, we will also be able to explore outcomes based on different age groups which is important because younger ages, for example, have a higher chance of the hip restoring to relative normality.

Acknowledgements. A. K. Davison was funded by Arthritis Research UK as part of the ORCHiD project. C. Lindner was funded by the Engineering and Physical Sciences Research Council, UK (EP/M012611/1) and by the Medical Research Council, UK (MR/S00405X/1). Manual landmark annotations were provided by the Medical Student Annotation Collaborative (Grace Airey, Evan Araia, Aishwarya Avula, Emily Gargan, Mihika Joshi, Muhammad Khan, Kantida Koysombat, Jason Lee, Sophie Munday, and Allen Roby).

References

1. Perry, D.: The epidemiology and etiology of Perthes’ disease. In: Koo, K., et al. (eds.) *Osteonecrosis*, pp. 419–425. Springer (2014), https://doi.org/10.1007/978-3-642-35767-1_58
2. Hall, A., Barker, D.: The age distribution of Legg-Perthes disease: an analysis using Sartwell’s incubation period model. *Am. J. Epidemiol.* 120(4), 531–536 (1984)
3. Wiig, O., Terjesen, T., Svenningsen, S., Lie, S.: The epidemiology and aetiology of Perthes disease in Norway: a nationwide study of 425 patients. *J. Bone Joint Surg. Br.* 88(9), 1217–1223 (2006), <https://doi.org/10.1302/0301-620X.88B9.17400>
4. Hunter, J.: (iv) Legg Calvé Perthes’ disease. *Curr. Orthop.* 18(4), 273–283 (2004), <https://doi.org/10.1016/j.cuor.2004.06.001>
5. Waldenstrom, H.: The first stages of coxa plana. *Acta Orthop. Scand.* 5(1–4), 1–34 (1934)
6. Joseph, B.: Natural history of early onset and late-onset Legg-Calve-Perthes disease. *J. Pediatr. Orthop.* 31(Suppl. 2), S152–S155 (2011), <https://doi.org/10.1097/BPO.0b013e318223b423>
7. Catterall, A.: Natural history, classification, and x-ray signs in Legg-Calvé-Perthes’ disease. *Acta Orthop. Belg.* 46(4), 346–351 (1980)
8. Salter, R., Thompson, G.: Legg-Calvé-Perthes disease: the prognostic significance of the subchondral fracture and a two-group classification of the femoral head involvement. *J. Bone Joint Surg. Am.* 66(4), 479–489 (1984)
9. Herring, J., Neustadt, J., Williams, J., Early, J., Browne, R.: The lateral pillar classification of Legg-Calvé-Perthes disease. *J. Pediatr. Orthop.* 12(2), 143–150 (1992)
10. Stulberg, S., Cooperman, D., Wallensten, R.: The natural history of Legg-Calvé-Perthes disease. *J. Bone Joint Surg. Am.* 763(7), 1095–1108 (1981)
11. Mouravliansky, N., Matsopoulos, G., Nikita, K., Uzunoglu, N., Pistevos, G.: An image processing technique for the quantitative analysis of hip disorder in Perthes’ disease. In: *Proc. 18th Annual International Conference of the IEEE Engineering in Medicine and Biology Society – EMBC 1996*, vol. 3, pp. 1103–1104. IEEE (1996), <https://doi.org/10.1109/IEMBS.1996.652728>
12. Chan, E., Farnsworth, C., Koziol, J., Hosalkar, H., Sah, R.: Statistical shape modeling of proximal femoral shape deformities in Legg-Calvé-Perthes disease and slipped capital femoral epiphysis. *Osteoarthr. Cartil.* 31(3), 443–449 (2013), <https://doi.org/10.1016/j.joca.2012.12.007>

13. Cootes, T., Taylor, C., Cooper, D., Graham, J.: Active shape models - their training and application. *Comput. Vis. Image Understand.* 61(1), 38–59 (1995), <https://doi.org/10.1006/cviu.1995.1004>
14. Cootes, T., Edwards, G., Taylor, C.: Active appearance models. *IEEE Trans. Pattern Anal. Mach. Intell.* 23(6), 681–685 (2001), <https://doi.org/10.1109/34.927467>
15. Breiman, L.: Random forests. *Mach. Learn.* 45(1), 5–32 (2001), <https://doi.org/10.1023/A:1010933404324>
16. Lindner, C., Thiagarajah, S., Wilkinson, J., The arcOGEN Consortium, Wallis, G., Cootes, T.: Fully automatic segmentation of the proximal femur using random forest regression voting. *IEEE Trans. Med. Imaging* 32(8), 1462–1472 (2013), <https://doi.org/10.1109/TMI.2013.2258030>
17. Lindner, C., Bromiley, P., Ionita, M., Cootes, T.: Robust and accurate shape model matching using random forest regression-voting. *IEEE Trans. Pattern Anal. Mach. Intell.* 37(9), 1862–1874 (2015), <https://doi.org/10.1109/TPAMI.2014.2382106>
18. Sarkalkan, N., Weinans, H., Zadpoor, A.: Statistical shape and appearance models of bones. *Bone* 60, 129–140 (2014), <https://doi.org/10.1016/j.bone.2013.12.006>
19. Waarsing, J., Rozendaal, R., Verhaar, J., Bierma-Zeinstra, S., Weinans, H.: A statistical model of shape and density of the proximal femur in relation to radiological and clinical OA of the hip. *Osteoarthr. Cartil.* 18(6), 787–794 (2010), <https://doi.org/10.1016/j.joca.2010.02.003>
20. Whitmarsh, T., Fritscher, K., Humbert, L., Del Rio Barquero, L., Roth, T., Kammerlander, C., Blauth, M., Schubert, R., Frangi, A.: A statistical model of shape and bone mineral density distribution of the proximal femur for fracture risk assessment. In: Fichtinger, G., et al. (eds.) *Proc. 14th International Conference on Medical Image Computing and Computer Assisted Intervention – MICCAI 2011*, *Lect. Notes Comput. Sc.*, vol. 6892, pp. 393–400. Springer (2011), https://doi.org/10.1007/978-3-642-23629-7_48
21. Thomson, J., O'Neill, T., Felson, D., Cootes, T.: Automated shape and texture analysis for detection of osteoarthritis from radiographs of the knee. In: Navab, N., et al. (eds.) *Proc. 18th International Conference on Medical Image Computing and Computer Assisted Intervention – MICCAI 2015*, *Lect. Notes Comput. Sc.*, vol. 9350, pp. 127–134. Springer (2015), https://doi.org/10.1007/978-3-319-24571-3_16
22. Adeshina, S., Cootes, T., Adams, J.: Automatic assessment of bone age using statistical models of appearance and random forest regression voting. In: *Proc. 13th International Conference on Electronics, Computer and Computation – ICECCO*, pp. 1–6. IEEE (2017), <https://doi.org/10.1109/ICECCO.2017.8333314>